

07-24-00

CERTIFICATE OF MAILING BY "EXPRESS MAIL"
(NEW PATENT APPLICATION)

Express Mail No. EL625291946US

Deposited July 21, 2000

I hereby certify that the attached correspondence, identified below, is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" under 37 CFR § 1.10 on the date indicated above and is addressed to the Asst. Commissioner for Patents, Box Patent Application, Washington, DC 20231.

By: B. Freeman
(person actually depositing)

In the Application of: Cloonan

Serial No.: _____

Filing Date: _____

Title: Congestion Control in a Network Device Having a Buffer Circuit

- ☒ Patent Application (23 pages, 22 claims, 5 sheets informal drawings)
- ☒ Return Postcard
- ☒ Transmittal Letter under 37 CFR 1.53(b) and duplicate
- ☒ Check in the amount of \$520.00
- ☒ Small Entity Form
- ☒ Declaration and Power of Attorney
- ☒ Assignment and Recordation Form Cover Sheet

Attorney Docket No.: 4807.00009

NEW UNITED STATES UTILITY PATENT APPLICATION
under 37 C.F.R. 1.53(b)

Page 1

Atty. Docket No. 4807.00009

CERTIFICATE OF EXPRESS MAILING UNDER 37 C.F.R. § 1.10: The undersigned hereby certifies that this United States Patent Application and all papers noted herein as being attached, are being deposited with the United States Postal Service "Express Mail Post Office to Addressee" Service under 37 C.F.R. § 1.10 today, **July 21, 2000**, and is addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231.

Express Mail Label No.: EL625291946US

Assistant Commissioner for Patents
Box Patent Applications
Washington, D.C. 20231

Enclosed herewith is a new patent application and the following papers:

First Named Inventor (or application identifier): Cloonan, Thomas J.

Title of Invention: Congestion Control in a Network Device Having a Buffer Circuit

1. ☒ Specification 23 pages (including specification, claims, abstract) / 22 claims (6 independent)
2. ☒ Declaration/Power of Attorney is:
 - ☒ attached in the regular manner.
 - ☐ NOT included, but deferred under 37 C.F.R. § 1.53(f).
3. ☒ 5 Distinct sheets of ☐ Formal ☒ Informal Drawings
4. ☐ Preliminary Amendment.
5. ☐ Information Disclosure Statement
 - ☐ Form 1449
 - ☐ A copy of each cited prior art reference
6. ☒ Assignment with Cover Sheet.
7. ☐ Priority is hereby claimed under 35 U.S.C. § 119 based upon the following application(s):

Country	Application Number	Date of Filing (day, month, year)

8. ☐ Priority document(s).

NEW UNITED STATES UTILITY PATENT APPLICATION
under 37 C.F.R. 1.53(b)

Page 2

Atty. Docket No. 4807.00009

9. ☒ Statement Claiming Small Entity Status.
10. ☐ Microfiche Computer Program (Appendix).
11. ☐ Nucleotide and/or Amino Acid Sequence Submission.
☐ Computer Readable Copy.
☐ Paper Copy (identical to computer copy).
☐ Statement verifying identity of above copies.

4807.00009

NEW UNITED STATES UTILITY PATENT APPLICATION
under 37 C.F.R. 1.53(b)

Page 3

Atty. Docket No. 4807.00009

12. Calculation of Fees:

FEES FOR	EXCESS CLAIMS	FEE	AMOUNT DUE
Basic Filing Fee (37 C.F.R. § 1.16(a))			\$690.00
Total Claims in Excess of 20 (37 C.F.R. § 1.16(c))	2	18.00	\$36.00
Independent Claims in Excess of 3 (37 C.F.R. § 1.16(b))	3	78.00	\$234.00
Multiple Dependent Claims (37 C.F.R. § 1.16(d))	0	270.00	\$0.00
Subtotal - Filing Fee Due			\$960.00
	MULTIPLY BY		
Reduction by 50%, if Small Entity (37 C.F.R. §§ 1.9, 1.27, 1.28)	0.5		\$480.00
TOTAL FILING FEE DUE			\$480.00
Assignment Recordation Fee (if applicable) (37 C.F.R. § 1.21(h))	1	40.00	\$40.00
GRAND TOTAL DUE			\$520.00

13. PAYMENT is:

- ☒ included in the amount of the GRAND TOTAL by our enclosed check. A general authorization under 37 C.F.R. § 1.25(b), second sentence, is hereby given to credit or debit our Deposit Account No. 01-0850 for the instant filing and for any other fees during the pendency of this application under 37 C.F.R. §§ 1.16, 1.17 and 1.18.
- ☐ not included, but deferred under 37 C.F.R. § 1.53(f).

14. All correspondence for the attached application should be directed to:

Banner & Witcoff, Ltd.
Ten South Wacker Drive, Suite 3000
Chicago, IL 60606-7407
Telephone: (312) 715-1000
Facsimile: (312) 715-1234

15. Other: _____

Date: July 21, 2000

By: Joseph P. Krause
Joseph P. Krause
Reg. No. 32,578

JPK/caj

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

(Attorney's Docket No. 4807.00009)

Applicant or

Patentee: Thomas J. Cloonan

Serial or
Patent No.

Filed or
Issued: _____

Title: Congestion Control in a Network Device Having a Buffer Circuit

**VERIFIED STATEMENT CLAIMING SMALL ENTITY STATUS
(37 C.F.R. § 1.9(f) AND § 1.27(c)) - SMALL BUSINESS CONCERN**

I hereby declare that I am

- ☐ the owner of the small business concern identified below:
☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF CONCERN Cadant, Inc.

ADDRESS OF CONCERN 4343 Commerce Ct., Suite 207

Lisle, Illinois 60532

I hereby declare that the above-identified small business concern qualifies as a small business concern as defined in 13 C.F.R. § 121.12, and reproduced in 37 C.F.R. § 1.9(d), for purposes of paying reduced fees to the United States Patent and Trademark Office, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time, or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention, entitled Congestion Control in a Network Device Having a Buffer Circuit

by inventor(s) Thomas J. Cloonan
described in _____

- ☒ the specification filed herewith
☐ Application Serial No. , filed .
☐ Patent No. , issued .

If the rights held by the above identified small business concern are not exclusive, each individual concern or organization having rights in the invention must file verified statements averring to their status as small entities, and no rights to the invention are held by any person, other than the inventor, who would not qualify as an independent inventor under 37 CFR § 1.9(c) if that person made the invention, or by any concern which would not qualify as a small business concern under 37 CFR § 1.9(d), or a nonprofit organization under 37 CFR § 1.9(e).

Each person, concern or organization having any rights to the invention is listed below:

☐ No such person, concern or organization exists.

☒ Each such person, concern or organization is listed below.

FULL NAME Cadant, Inc.

ADDRESS 4343 Commerce Ct., Suite 207, Lisle, Illinois 60532

☐ Individual ☒ Small Business Concern ☐ Nonprofit Organization

FULL NAME _____

ADDRESS _____

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

Separate verified statements are required from each named person, concern or organization having rights in the invention averring to their status as small entities. (37 CFR § 1.27).

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 C.F.R. § 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing therein, or any patent to which this verified statement is directed.

Venkata C. Majeti


NAME OF PERSON SIGNING

President and CEO

TITLE IN ORGANIZATION

4343 Commerce Court, Suite 207, Lisle, IL 60532

ADDRESS OF PERSON SIGNING



Signature

July 19, 2000

Date

CONGESTION CONTROL IN A NETWORK DEVICE HAVING A BUFFER CIRCUIT

FIELD OF THE INVENTION

The present invention relates generally to data flow control. Particularly, the present
5 invention relates to control of congestion in a buffer circuit.

DESCRIPTION OF THE RELATED ART

In order to provide more service to their subscriber base, cable television companies are
offering access to the Internet through their cable modem (CM) boxes. The benefits in using the
cable companies instead of a dial-up Internet Service Provider is multiple services under one bill,
10 always-on access, and, in some cases, higher speed access.

In order to provide their customer's with Internet access, the cable companies use some
of the 50 – 800 MHz spectrum typically set aside for their television channels to provide the
bandwidth required for the data transfers. A typical cable system has the bandwidth to provide
100 television channels to its subscribers. Each NTSC television signal requires 6 MHz of
15 bandwidth.

In order for a cable subscriber to access the Internet through their cable television
provider, the subscriber must have a CM. The CM is similar to the Cable Modem Termination
System (CMTS) equipment required at the cable company's headquarters, except for the greater
size required at the headquarters. This is to accommodate a greater number of signals than is
20 required by the home modem.

The home CM box and the CMTS use well-known Ethernet frames to communicate
between them. The cable system, however, uses a different modulation scheme, Quadrature
Amplitude Modulation (QAM), than is normally used in an Ethernet scheme.

Using the QAM modulation, the downstream (from the cable company equipment to the
25 home CM) data rate is in the range of 30-40 Mbps for each 6 MHz channel. This can

accommodate between 500 and 2000 subscribers. The more subscribers that the cable company tries to fit in that spectrum, the lower the bandwidth available for each subscriber.

The upstream data flow is different and more complex. In the past, cable companies did not have to worry about providing bandwidth for the customer to communicate in the upstream
5 direction. Pay-per-view movies and sports events, however, required this ability. The cable companies, therefore, set aside the 5 – 42 MHz spectrum to allow the home CM to communicate in the upstream direction. The cable companies now use this 5 – 42 MHz spectrum to provide the necessary upstream access to the Internet from the home CM.

Cable companies, as well as other Internet Service Providers, are currently introducing
10 Quality of Service (QoS) to Internet access. The current Internet routing model of “best effort” service now provided to all users, packets, and traffic flows is being replaced with services that differentiate between packets.

Some current Internet access schemes to increase QoS incorporate different priority
levels for data packets. High priority packets will be routed with low latency and low jitter while
15 low priority packets may experience more delay and jitter. The throughput needs of each application determine the priority associated with its corresponding traffic flows. It is likely that advanced application programs in the future will dynamically change the priority of traffic flows to match the varying needs of the user throughout the entire duration of the user’s session.

Another method presently used by cable companies to introduce higher quality Internet
20 access service is to change the techniques presently used in dropping all overflow packets, increasing data throughput, and decreasing data throughput. If packets are dropped in a synchronized fashion, this synchronization of drops can lead to an inefficient use of the cable system bandwidth. Attempts at de-synchronization include the Random Early Discard (RED) scheme, which randomly drops user packets based on the current buffer depth within storage
25 buffers in the cable data system.

FIG. 1 illustrates a typical prior art CMTS block diagram. The CMTS typically is comprised of a cable interface card (101) to provide the interface signals and modulation to the signals transmitted to the home modem. An Ethernet card (110) interfaces the CMTS to the Internet by providing appropriate timing, control, and data signal formats for the Internet. A
5 buffer circuit (105) between the cable interface card (101) and Ethernet card (110) stores data in both the upstream and downstream directions when the processing in either the cable interface card or the Ethernet card is slower than the incoming data.

The Random Early Discard scheme samples the depth (amount of memory space used) of the buffer (105) and randomly drops packets to prevent the buffer from overflowing. Using this
10 scheme, a cable subscriber may not lose any data or may lose a lot of data, depending on a purely random occurrence.

In the future, since all packets will not be passed using the same priority levels, all packets will not be billed the same charges. A subscriber who is paying more for his service will expect higher data throughput and less data loss than someone paying substantially less. There is
15 a resulting need for an improved process and system for determining which data packets to drop before the buffer becomes overloaded. One technique that performs this task is the Weighted Random Early Discard (WRED) scheme, which randomly drops user packets based on current buffer depths and based on the priority of the packet. Packets associated with high-paying customers will be given higher priority and experience a lower probability of packet drops than
20 packets associated with low-paying customers.

Although the use of the Weighted Random Early Discard scheme serves to differentiate users based on their priorities (which is usually related to their payment), it does not differentiate between two different customers who are assigned to the same priority level. In particular, the Weighted Random Early Discard scheme does not take into account the customer's recent
25 utilization of the bandwidth. As a result, two customers that are both assigned to the same

priority might be utilizing the bandwidth in two very different ways. One might have been actively using the bandwidth a lot in the last few minutes, and the other might not have been using the bandwidth at all during the last few minutes. Because of his or her low bandwidth activity levels in recent minutes, it seems fair that the second customer should be given

5 precedence over the “recently greedy” first customer if both customers start to transmit packets simultaneously. The ability to store the activity state of each customer and use that activity state in the calculation of packet dropping probabilities in addition to the priority and the current packet buffer depths would therefore seem to be a useful and fair manner of controlling the flow of packets in a data communications system.

10 SUMMARY OF THE INVENTION

The present invention encompasses a process for providing data packet flow congestion control in a data network that has a buffer circuit. Each packet of the present invention is identified as belonging to a specific service flow, where a service flow is a series of packets that share something in common. The shared attributes could, for example, include a common source
15 IP address and a common destination IP address within the packet headers of all of the packets within the service flow. A packet stream will typically contain many packets from many different service flows inter-mixed together. When each of these packets arrives at a network node, it must then be assigned a priority level. In the preferred embodiment, the priority level is based on the price a subscriber paid for their Internet access service.

20 The process begins by detecting the current data packet rate through the buffer circuit for the service flow associated with the arriving packet. This service flow’s data packet rate is quantized into at least one activity level. In the preferred embodiment, the process uses four different activity levels for data rate. Each level is determined by a comparison to predetermined thresholds.

The current buffer circuit depth is then determined as is the priority associated with the current data packet. The current packet is either dropped or used depending on the current data packet flow rate, the data packet priority, and the current buffer circuit depth.

5

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a block diagram of a typical prior art cable modem termination system.

FIG. 2 shows a block diagram of the cable modem termination system apparatus of the present invention.

FIG. 3 shows a matrix of settings for congestion control on the Cable Interface circuit
10 and the Ethernet Interface circuit.

FIG. 4 shows a plot of the probability of dropping a packet versus the average buffer depth in accordance with the matrix of FIG. 3.

FIG. 5 shows a flowchart of the buffer congestion control process of the present invention.

FIG. 6 shows a block diagram of a system incorporating the cable modem apparatus of
15 the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides congestion control of any data network circuit having a buffer. By tracking the data packet flow rate, data packet priority and buffer status, the packets
20 can be dropped before they overflow the buffer.

FIG. 2 illustrates the preferred embodiment cable modem termination system (CMTS) apparatus of the present invention. While the present invention will be described as applying to a cable modem, the congestion control can be applied to any data network comprising a buffer circuit.

The CMTS apparatus of FIG. 2 is comprised of a cable interface (201) that is coupled to a buffer circuit (205). The buffer circuit (205) is coupled to an Ethernet interface (210). In the preferred embodiment, each of the individual circuits (201, 205, and 210) reside physically on separate circuit boards. In alternate embodiments, any circuits having substantially the same function can reside on one circuit board or even one integrated circuit. In other words, the present invention is not limited to three separate circuit boards.

The cable interface (201) is responsible for interfacing the CMTS to the home cable modem apparatus. The cable interface (201) also provides the functions of modulation and demodulation.

The cable interface circuit is comprised of a downstream packet flow path and an upstream packet flow path. The downstream packet flow path is comprised of a data throughput monitor (220) that is coupled to a flow limiter (215). The data throughput monitor (220) has an input that is coupled to the buffer circuit (205) from which the data packets flow and a feedback from the upstream path. The feedback from the upstream path is to allow a first CM to talk with other CMs. The data throughput monitor (220) has the task of determining the rate of data packet flow.

In the preferred embodiment of the CMTS, the downstream data packet flow rate is typically either 30 or 40 Mbps for each 6 MHz channel, using QAM techniques. Alternate embodiments use other flow rates. The cable company decides which data packet flow rate depending on the outcome desired by the company. The lower data rate is less susceptible to noise while the higher data rate can include more data per unit of time for the customers.

The data packet flow rate signal is fed into the flow limiter (215). This signal controls the flow limiter function. If the flow is greater than a predetermined level, T_{\max} , the data packet flow can be limited. The flow limiter (215) reduces the data rate by dropping packets until the flow is reduced to below T_{\max} .

Another input to the flow limiter (215) is the “limiting type” input. This control input is set by the cable company depending on how strict they wish a customer to adhere to the rules. If the “limiting type” input is set to “soft-limiting”, the flow limiter (215) allows the data rate to go above the set data rate by a predetermined amount without dropping any packets.

5 Some cable companies may strictly limit a customer to T_{\max} . In this case, the “limiting type” control input is set to “hard-limiting”. If the data rate goes over the set hard limit, the flow limiter (215) drops any packets that force the customer to exceed T_{\max} . The output of the flow limiter (215) is coupled to the cable that runs to the customers’ cable modems.

10 The output of the flow limiter (215) is input to the modulator (255). This block (255) performs the QAM needed to transmit the data to the CMs.

15 The upstream data path is comprised of a demodulator and filter (260) that converts the QAM signal into data bits in order to be processed by the other blocks in the upstream path. The demodulated data bits are input to a data throughput monitor (225) that is coupled to the upstream port from the customer’s CM. This data throughput monitor (225) has the same functionality as the downstream monitor (220) of monitoring the data rate but in the upstream direction to the Internet.

In the preferred embodiment, the upstream data rate can be in the range of 320 kb to 10.24 Mbps. Alternate embodiments use other rates.

20 The upstream data throughput monitor (225) is coupled to a flow limiter (230). This flow limiter has similar functionality to the flow limiter (215) in the downstream path. The upstream path flow limiter (230) has the data rate input from the data throughput monitor (225) as well as the “limiting type” control input that, in the preferred embodiment, is set to either “hard-limiting” or “soft-limiting” depending on the cable company rules. As in the downstream flow limiter (215), the upstream flow limiter, depending on the “limiting type” input, drops all packets
25 that force the customer to exceed T_{\max} .

The upstream path further comprises a congestion control block (235) that is coupled to the upstream data path out of the flow limiter (230). The data packets from the upstream data path flow through the congestion control block (235) to the buffer circuit (205). The function of the congestion control block (235) is to drop packets when the buffer depth is reaching a
5 maximum point. By dropping the packets before they reach the buffer, the buffer will not overflow.

In order to accomplish the task of congestion control, the congestion control block (235) has control inputs that are used to determine when to drop packets and which packets to drop. In the preferred embodiment, these control inputs include the data rate signal from the upstream
10 data throughput monitor (225), a buffer depth signal from the buffer (205), and a priority signal.

The data rate signal from the upstream data throughput monitor (225), as described above, quantizes the data rate and feeds that value to the congestion control block (235). The buffer circuit depth signal from the buffer circuit (205) instructs the congestion control block (235) as to the depth of the buffer. In other words, if the buffer (205) is 75% full, the buffer
15 depth signal instructs the congestion control block (235) of this.

The priority signal that is input to the congestion control block (235) informs the congestion control of the priority of each packet. This is important in determining which packets to drop.

A group of packets is assigned a priority based on the customer's level of service plan. If
20 the customer has signed up for the basic service plan and paid the smallest fee for the most basic service, his packets are assigned a low priority. This priority is embedded in a packet identification that is assigned to the group of packets and is decoded when the group of packets enters the cable interface.

If the customer has signed up for the premium service plan with the cable company, his
25 packets are assigned the highest priority. If the customer has signed up for any service plans that

are in between the premium and the basic plans, this priority is also assigned to each packet. As described before, the priority is added to the packet identification for a particular group of packets.

A customer may also decide to dynamically change his service level for a given session.

- 5 In this case, different packet groups from that particular customer will have different priorities assigned to different packet identifications.

As described subsequently in other figures, the congestion control block (235) of FIG. 2 uses the priority assigned to a group of packets to determine how to process that particular group of packets. The output of the congestion control block is input to the buffer circuit's upstream
10 data flow input.

The buffer circuit (205) stores the packets until the Ethernet circuit (210) has time to process that packet. The packets are fed from the buffer circuit (205) to the Ethernet circuit (210) as more processing time is freed up.

The downstream path of the Ethernet circuit (210) is comprised of a data throughput
15 monitor (250) that is coupled to the connection to the Internet. This monitor (250) provides substantially the same function as the previously described data throughput monitors on both the upstream and downstream paths.

The data packets from the Internet flow from the data throughput monitor (250) to the Ethernet's circuit flow limiter (245). This flow limiter (245) has substantially the same
20 functionality as the above described flow limiters. This flow limiter also has the same inputs as described previously: the quantized data rate and the "limiting type" control input.

The data packets flow from the flow limiter (245) to the congestion control block (240). As in the upstream congestion control block (235), the Ethernet's downstream congestion control block (240) has the three control inputs to determine which packets to drop: the quantized data

rate, the buffer depth signal, and the packet priority signal. The congestion control block then drops a particular packet based on these control signals.

The downstream data flows from the congestion control block to the buffer circuit (205). The buffer circuit (205) stores the packets until the cable interface circuit has the processing time to work on additional packets.

The buffer circuit (205) is comprised of 128 MB of RAM, in the preferred embodiment. Alternate embodiments use other values of RAM or even other types of memory instead of RAM. The alternate types of memory include hard drives or other types of temporary memory.

Most of the functions illustrated in FIG. 2 may be implemented in various ways. These functions can be performed in software by a processor or multiple processors performing each function. Each function can also be implemented in discrete logic hardware, a digital signal processor, or some other form of programmable logic.

FIG. 3 illustrates a matrix of settings for congestion control of the buffer circuit. These settings are an example of the preferred embodiment settings that were determined experimentally in field tests. Alternate embodiments use other settings.

The matrix illustrates the priority and the data flow rate. The priority indicates the bandwidth for which the customer has paid. The data flow rate indicates the bandwidth that the customer is actually using.

The matrix of FIG. 3 is comprised of columns under the current data flow rate across the top of the matrix. The data rate, in the preferred embodiment, is segmented into four different rates that are labeled needy, normal, greedy, and super-greedy. These data rates are based on the minimum bandwidth for which the customer signed up (T_{\min}) and the maximum data rate that a customer requested during a session (T_{\max}).

The data rate is considered “needy” when the rate is between 0 and T_{\min} . If the data rate is between T_{\min} and T_{\max} , this rate is considered normal. When the data rate goes above some mid-

rate (T_{mid}) and T_{max} , this rate is labeled greedy. When a customer requests a data rate beyond T_{max} , the rate is considered super-greedy.

These labels and respective data rates are for illustration of the operation of the preferred embodiment only. Alternate embodiments use other labels for other data rate thresholds.

Referring again to the matrix of FIG. 3, the rows are labeled with six different priority levels that have been assigned to each data packet as the Current Packet Priority. The highest priority level is designated Priority 0. The lowest priority level is designated Priority 5. The mid-level priorities are labeled Priority 1, Priority 2, Priority 3, and Priority 4 in increasing priority levels.

The blocks that make up the intersection of each column (current data flow rate) and row (current packet priority) is comprised of three values that are further illustrated in the plot of FIG. 4 to be discussed subsequently. The upper number in each block represents the minimum threshold value (in percent) of the average buffer depth (amount of memory used). This value is subsequently referred to as MIN_{th} . The middle value in each block is the maximum threshold (in percent) of the average buffer depth. This value is subsequently referred to as MAX_{th} .

In the preferred embodiment, MIN_{th} and MAX_{th} are average values. Alternate embodiments may use instantaneous values of the buffer depth for these thresholds.

The third value in each block is the probability of a packet being dropped. In the illustrated example, the probability of a packet being dropped is 0.1 and is referred to as P_a . This value for P_a is a typical value and will be different for other embodiments.

The “needy” data rate column ensures that T_{min} is guaranteed. The “super-greedy” column ensures that T_{max} is a pseudo-hard limit. This limit is set by the cable company and may or may not be adhered to in limiting a customer’s request for additional bandwidth. This is mostly a function of the buffer depth. If the buffer depth is low and customers are above their paid bandwidth limit, their packets will not be dropped. However, once the buffer approaches its

limit and customers are above their paid bandwidth limit, their packets will be dropped to keep them from going over T_{\max} .

The values of the matrix of FIG. 3 were determined by expanding on the concepts in Weighted Random Early Discard (WRED). WRED is similar to the RED technique described above except it assigns a high priority (lower packet loss probability) to some data packets and a low priority (higher packet loss probability) to other data packets. For example, voice may be assigned a high priority while email may be assigned a low priority. In a congestion situation, the lower priority packets are discarded before the higher priority packets.

If WRED were used to determine the values in the blocks of the FIG. 3 matrix, all of the entries across a given row would be the same. If RED were used to determine the values, all of the entries would be the same. Using an extension of WRED, different values can be determined for each row of the matrix.

Each block of the matrix of FIG. 3 determines the appearance of the plot illustrated in FIG. 4. The plot of FIG. 4 has the average buffer depth running along the x-axis and the probability of dropping a packet (P_{drop}) running along the y-axis.

This plot illustrates that as the average buffer depth approaches MIN_{th} , the probability of a packet being dropped is 0%. This is due to the fact that the buffer depth is not yet high, so there is no need to drop any packets.

As the average buffer depth increases along the x-axis and approaches MAX_{th} , the probability of a packet being dropped increases to P_a . A typical value of P_a , as illustrated in the matrix of FIG. 3, is 0.1. When the average buffer depth exceeds MAX_{th} , all packets are dropped with a probability of 1.0.

As a safety mechanism to keep the buffer circuit from overflowing, once the instantaneous buffer depth goes beyond 99%, all packets are dropped. The priority of the packet is irrelevant.

The process used by the present invention is illustrated in FIG. 5. The process begins with the CM registering with the cable system according to the Data Over Cable Service Interface Standard (DOCSIS).

5 The CM and the CMTS go through a ranging process (step 501) to get the CM on a channel that is acceptable to the CMTS. The ranging step (step 501) also includes determining the power level required by the CMTS from the CM.

10 The CM next goes through the registration process (step 505) to get the information required to be an Internet Protocol (IP) host, such as the IP address, from the Dynamic Host Configuration Protocol (DHCP) server. During this registration process, the CM also downloads, from the a Trivial File Transfer Protocol (TFTP) server, the service level and other service information for that particular CM.

15 The process of the present invention enables the customer to change his service level temporarily and dynamically, depending on the bandwidth he needs at the moment. This can be accomplished by the customer informing the CMTS of the additional bandwidth request (step 510).

The CMTS then informs the CM of the new service level (step 515). As the customer is using the service, the CMTS is monitoring the data flow rate (step 520) so that it knows the amount of bandwidth the customer is using in relation to the bandwidth for which the customer has paid.

20 The data packet rate is quantized to one of four levels (step 525), depending on the rate required by the customer's use. These levels, as discussed above, include needy, normal, greedy, and super-greedy.

25 Since the current packet priority is in the header of each packet, the process determines the current packet priority (step 530) by reading the packet header. The process also determines the average buffer depth (step 535) to determine the amount of memory remaining in the buffer.

By using the data rate, current packet priority, and average buffer depth, the process accesses the proper block of the matrix of FIG. 3 to determine the probability that a packet should be dropped (step 540).

FIG. 6 illustrates a system block diagram of the present invention. This system is
5 comprised of the CMTS (601) that is coupled to the Internet (610). The CMTS is located at the cable company headquarters.

The CMTS is coupled to a larger number of CMs (612 – 614). For purposes of clarity, only three CMs are shown. However, as described above, the CMTS may communicate with thousands of CMs, depending on how the cable company set up the system.

10 In summary, the process and apparatus of the present invention provides congestion control of a buffer circuit. By tracking the depth of the buffer circuit, the data rate, and the priority of each packet, packets can be selectively dropped prior to overflowing the buffer. This also allows the customer to dynamically change his service level depending on his bandwidth needs for each Internet access session. By embedding the new priority in each data packet, each
15 data packet can be dropped or processed based on its own priority.

WE CLAIM:

1 1. A method for providing data packet congestion control for a data network having
2 a buffer circuit, each data packet comprising a priority, the method comprising the steps of:
3 determining the particular service flow associated with the data packet,
4 detecting a current data packet flow rate through the data network for the particular
5 service flow associated with the data packet;
6 quantizing the data packet flow rate into at least one level;
7 detecting a buffer circuit depth;
8 determining the priority associated with a current data packet; and
9 processing the current packet in response to the current data packet flow rate, the data
10 packet priority, and the current buffer circuit depth.

1 2. The method of claim 1 wherein the step of detecting a buffer circuit depth
2 includes detecting an instantaneous value for the buffer circuit depth.

1 3. The method of claim 1 wherein the step of detecting a buffer circuit depth
2 includes detecting an average value for the buffer circuit depth.

1 4. The method of claim 1 wherein the step of processing includes determining a
2 probability of dropping a data packet using the current data packet flow rate, the data packet
3 priority, and the current buffer circuit depth to access, from a matrix of plots indicating the
4 probability of dropping the data packet, each of the plots within the matrix of plots indicating the
5 packet drop probability as a function of the buffer circuit depth.

1 5. The method of claim 1 wherein the step of processing includes determining a
2 probability of dropping a data packet using the current data packet flow rate, the data packet
3 priority, and the current buffer circuit depth to access, from a matrix of plots indicating the
4 probability of dropping the data packet, each of the plots within the matrix of plots indicating the
5 packet drop probability as a function of the buffer circuit depth with a plots that have the same
6 shape and parameters as the plots used for Random Early Discard congestion control techniques.

1 6. The method of claim 1 wherein the step of quantizing includes quantizing the data
2 packet flow rate into four different flow rates.

1 7. The method of claim 6 wherein the four predetermined flow rates are determined
2 by comparing the service flow's data packet flow rate with a minimum data rate threshold, a
3 maximum data rate threshold, and a mid-level data rate threshold.

1 8. The method of claim 1 wherein the step of detecting the buffer circuit depth
2 includes determining the percentage of utilized memory within the buffer circuit.

1 9. The method of claim 1 wherein the step of determining the priority associated
2 with a current data packet includes the step of reading the header of the current data packet to
3 determine the priority.

1 10. The method of claim 1 wherein the step of determining the particular service flow
2 associated with the data packet includes the step of reading the header of the current data packet
3 to determine the associated service flow.

1 11. A method for adjusting the data rate of data packets transiting a data network, the
2 system comprising:

3 determining a maximum data rate of data packet flow allowed through the data network;

4 monitoring the data rate of the data packets through the data network;

5 determining a predetermined priority of a data packet;

6 determining remaining memory in a memory circuit;

7 if the data packet has a low priority and the data rate is greater than the maximum data
8 rate, dropping the data packet; and

9 if the data packet has a high priority and the data rate is greater than the maximum data
10 rate, processing the data packet.

1 12. The method of claim 11 and further including the steps of:

2 determining a minimum data rate of data packet flow allowed through the data network;

3 and

4 if the data rate is less than the minimum data rate, always processing the data packet.

1 13. An apparatus for controlling the congestion of a buffer circuit, having a depth
2 indication signal, in a data network, the apparatus comprising:

3 the data monitor coupled to the data network input and accepting a stream of data

4 packets, the data rate monitor outputting the stream of data packets and a control signal

5 indicating the service flow's data packet flow rate;

6 a flow limiter having a data input coupled to the stream of data packets output from the

7 data rate monitor, the flow limiter also having a control input coupled to the control signal from

8 the data rate monitor, the flow limiter outputting the stream of data packets; and

9 a congestion controller having a data input coupled to the stream of data packets from the
10 flow limiter, the congestion controller also having a first control input coupled to the control
11 signal from the data rate monitor, a second control input coupled to the depth indication signal,
12 and a third control input coupled to a priority signal, the congestion controller controlling the rate
13 at which the stream of data packets enters the buffer circuit in response to the first, second, and
14 third control signals by dropping predetermined data packets from the stream of data packets.

1 14. The apparatus of claim 13 wherein the flow limiter further includes a limiting
2 type control input for selectively changing an allowed data packet flow rate.

1 15. The apparatus of claim 13 wherein the data rate monitor is coupled to a plurality
2 of cable modems.

1 16. A cable modem interface for interfacing a cable modem termination system to a
2 plurality of cable modems, the cable modem termination system having a buffer circuit for
3 temporarily storing input data, the buffer circuit having a buffer depth indication signal that
4 indicates a depth of the buffer circuit, the interface comprising:

5 a downstream data path having a data stream input and a modulated output, the
6 downstream data path comprising:

7 a first data throughput monitor coupled to the data stream input, the first data
8 throughput monitor tracking a data rate of the data stream through the first data
9 throughput monitor and generating a first data rate signal in response to the tracked data
10 rate;

11 a first flow limiter, coupled to the first data throughput monitor, for limiting the
12 data rate of the data stream, the flow limiter having a control input coupled to the first
13 data rate signal; and

14 a modulator, coupled to the first flow limiter, for modulating the data stream for
15 transmission to the plurality of cable modems; and

16 an upstream data path comprising:

17 a demodulator, coupled to the plurality of cable modems, for demodulating an
18 input modulated signal into an input data bit stream comprising a plurality of packets;

19 a second data throughput monitor for monitoring a data flow rate of the input data
20 bit stream and generating a second data rate signal;

21 a second flow limiter, coupled to the second data throughput monitor, for limiting
22 the data flow rate of the input data bit stream, the second flow limiter coupled to the
23 second data rate signal;

24 a congestion controller, coupled to the second flow limiter, the congestion
25 controller having a plurality of control inputs, a first control input coupled to the buffer
26 depth indication signal, a second control input coupled to the second data rate signal, and
27 a third control input coupled to a priority signal indicating a priority of a group of
28 packets, the congestion controller controlling the number of packets entering the buffer
29 circuit in response to the plurality of control inputs.

1 17. The cable modem interface of claim 16 and further including a limiting type input
2 to the first flow limiter for determining adherence of the cable modem interface to a set of cable
3 modem rules.

1 18. The cable modem interface of claim 16 and further including a limiting type input
2 to the second flow limiter for determining adherence of the cable modem interface to a set of
3 cable modem rules.

1 19. A method for a cable modem to communicate with a cable modem termination
2 system, the cable modem being assigned a predetermined service level, the cable modem
3 termination system comprising a buffer circuit for temporarily storing data packets, the method
4 comprising the steps of:

5 the cable modem and the cable modem termination system performing a ranging step to
6 determine a channel and power level for communication that is acceptable to the cable modem
7 termination system;

8 the cable modem registering with the cable modem termination system;

9 the cable modem requesting additional bandwidth beyond the predetermined service
10 level;

11 communicating data packets from the cable modem to the cable modem termination
12 system;

13 determining a priority of the data packets;

14 determining an average depth of the buffer circuit;

15 determining a service flow associated with the data packets;

16 determining a flow rate of the service flow associated with the data packets; and

17 processing the data packets in response to the average depth of the buffer circuit, the flow
18 rate, and the priority.

1 20. The method of claim 19 wherein the step of processing includes dropping packets
2 in response to the average depth of the buffer circuit, the flow rate, and the priority.

1 21. The method of claim 19 wherein the step of registering includes the cable modem
2 termination system assigning an Internet protocol address to the cable modem.

1 22. A cable modem comprising:

2 a buffer circuit comprising:

3 a buffer depth indication signal that indicates the amount of memory not being
4 used;

5 a cable modem interface comprising:

6 a plurality of data rate monitors, a first monitor for monitoring a data packet flow
7 rate in an upstream path and a second monitor for monitoring the data packet flow rate in
8 a downstream path, each monitor generating a flow rate indication signal;

9 a plurality of flow limiters, a first flow limiter, coupled to the first data rate
10 monitor, limiting the flow rate in the upstream path and a second flow limiter, coupled to
11 the second data rate monitor, limiting the flow rate in the downstream path, both flow
12 limiters acting in response to the respective flow rate indication signal; and

13 a congestion controller, coupled to the first flow limiter, for controlling the flow
14 of data into the buffer circuit in response to the flow rate indication signal, the buffer
15 depth indication signal, and a packet priority signal generated from the data packets; and
16 an Ethernet interface comprising:

17 a data rate monitor, coupled to the Internet, for monitoring the data packet flow
18 rate in the downstream path, the data rate monitor generating a flow rate indication
19 signal;

20 a flow limiter, coupled to the data rate monitor, for limiting the flow rate in the
21 downstream path, the flow limiter acting in response to the flow rate indication signal;
22 and
23 a congestion controller, coupled to the flow limiter, for controlling the data
24 packets flowing into the buffer circuit in response to the flow rate indication signal, the
25 buffer depth indication signal, and a packet priority signal generated from the data
26 packets.

ABSTRACT OF THE DISCLOSURE

Each packet of the present invention is assigned a priority level. The current data packet flow rate is detected. This data packet flow rate is quantized into at least one data rate level. The current buffer circuit depth is determined as is the priority associated with the current data
5 packet. The probability that the current packet is either dropped or used is determined by using the current data packet service flow rate, the data packet priority, and the current buffer circuit depth.




FIG. 1

- PRIOR ART -

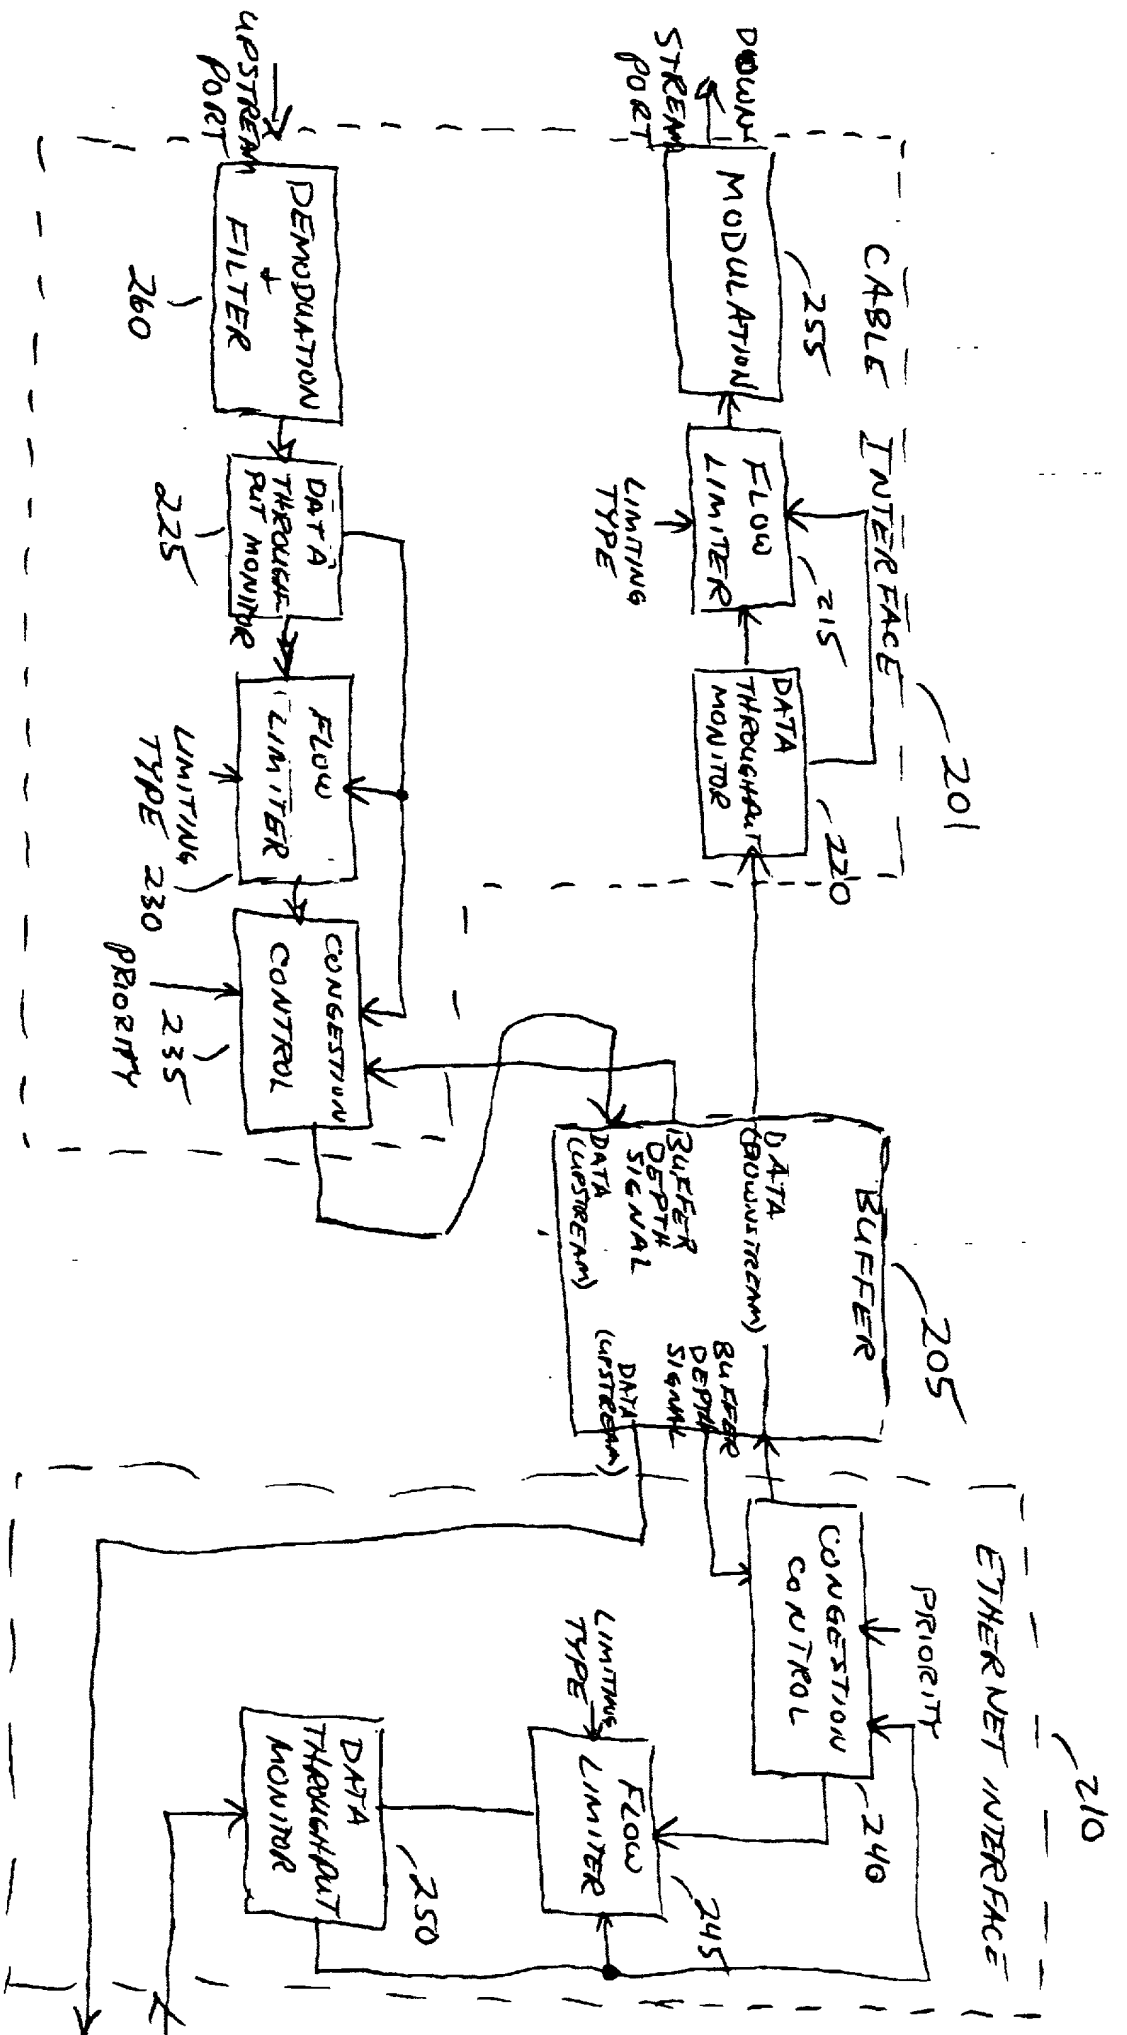


FIG. 2


FIG. 3

Current Packet Priority	MINth MAXth Pa	CURRENT FLOW RATE			
		Needy	Normal	Greedy	Super-Greedy
Priority 0		98% 99% 0.1	75% 95% 0.1	65% 85% 0.1	5% 25% 0.1
Priority 1		98% 99% 0.1	65% 85% 0.1	55% 75% 0.1	5% 25% 0.1
Priority 2		98% 99% 0.1	55% 75% 0.1	45% 65% 0.1	5% 25% 0.1
Priority 3		98% 99% 0.1	45% 65% 0.1	35% 55% 0.1	5% 25% 0.1
Priority 4		98% 99% 0.1	35% 55% 0.1	25% 45% 0.1	5% 25% 0.1
Priority 5		98% 99% 0.1	25% 45% 0.1	15% 35% 0.1	5% 25% 0.1

ensures T_{min} is guaranteed

ensures T_{max} is a pseudo-hard limit

FIG. 4



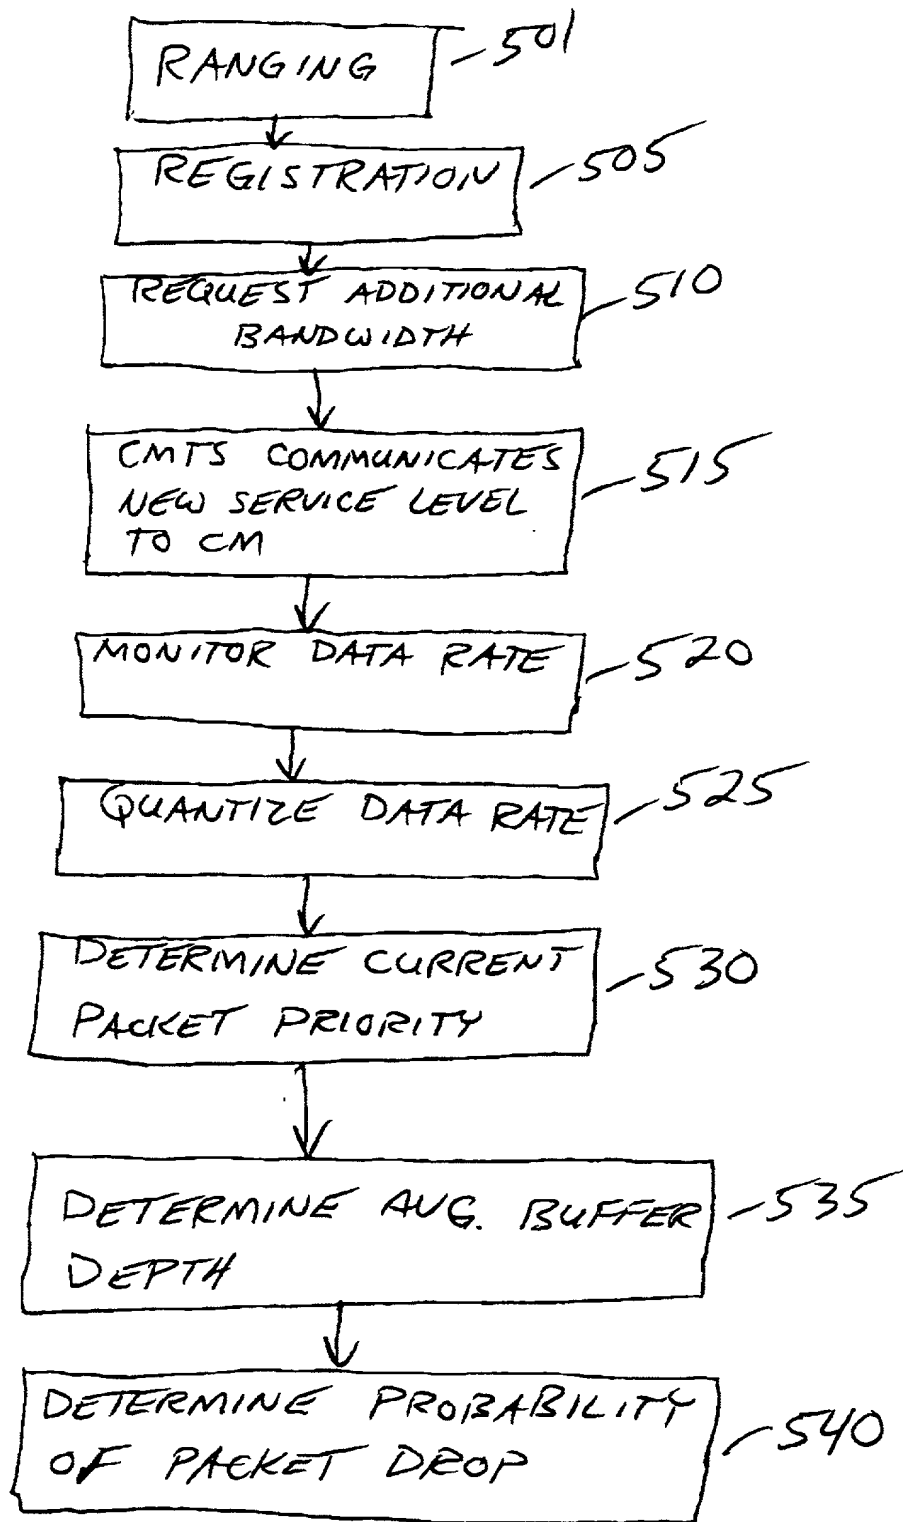



FIG. 5



DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Congestion Control in a Network Device Having a Buffer Circuit

the specification of which is attached hereto unless the following space is checked:

_____ was filed on _____ as United States Application Serial Number or PCT International Application Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR § 1.56.

I hereby claim foreign priority benefits under 35 U.S.C. § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT international application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s):

	<u>Number</u>	<u>Country</u>	<u>Day/Month/Year Filed</u>	<u>Priority Not Claimed</u>
1.				9
2.				9

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

	<u>Application Number</u>	<u>Filing Date</u>
1.		
2.		

I hereby claim the benefit under 35 U.S.C. § 120 of any United States application(s), or § 365(c) of any PCT international application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

	<u>Application Number</u>	<u>Filing Date</u>	<u>Status</u> <input checked="" type="checkbox"/> <u>patented</u> , <input type="checkbox"/> <u>pending</u> , <input type="checkbox"/> <u>abandoned</u>
1.			
2.			

I hereby appoint the following attorneys and agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

ALTHER, Robert F.	31,810	HYMEL, Lin J.	P45,414	PAYNE, Stephen S.	35,316
BANNER, Donald W. 17,037		IWANICKI, John P.	34,628	PETERSON, Thomas L.	30,969
BANNER, Mark T.	29,888	JACKSON, Thomas H.	29,808	POTENZA, Joseph M.	28,175
BANNER, Pamela I.	33,644	KAGAN, Sarah A.	32,141	PRATT, Thomas K.	37,210
BECKER, Matthew P. 45,824		KATZ, Robert S.	36,402	RENK, Christopher J.	33,761
BECKETT, William W.	18,262	KLEIN, William J.	43,719	RESIS, Robert H.	32,168
BERGHAMMER, Joseph J.	46,057	KRAUSE, Joseph P.	32,578	RIVARD, Paul M.	43,446
BODNER, Jordan	42,338	LINEK, Ernest V	29,822	SCHAD, Steven P.	32,550
BUROW, Scott A.	42,373	MAGOON, Sumeet	43,769	SHIFLEY, Charles W	28,042
CALLAHAN, James V.	20,095	MALONE, Dale A.	32,155	SKERPON, Joseph M.	29,864
CHANG, Steve S	42,402	MANNAVA, Ashok K.	45,301	SPAR, Elizabeth	45,123
COHAN, Gregory J.	40,959	MAPLE, Marie-Claire B.	37,588	STOCKLEY, D. J.	34,257
COOPERMAN, Marc S.	34,143	MAY, Steven A.	44,912	VAN ES, J. Pieter	37,746
CURTIN, Joseph P.	34,571	McDERMOTT, Peter D.	29,411	WILLIAMS, Kathleen M.	34,380
DAWSON, John R.	39,504	McKEE, Christopher L.	32,384	WITCOFF, Sheldon W.	17,399
DEMOOR, Laura J.	39,654	McKIE, Edward F	17,335	WOLFF, Kevin A.	42,233
EVANS, Thomas L.	35,805	MEDLOCK, Nina L.	29,673	WOLFFE, Franklin D. 19,724	
FEDOROCKO, Gary D.	35,509	MEECE, Timothy C.	38,553	WOLFFE, Susan A.	33,568
FISHER, Daniel E.	34,162	MEEKER, Frederic M.	35,282	WRIGHT, Bradley C.	38,061
FISHER, William J.	32,133	MILLER, Charles L.	43,805		
GLEBOCKI, Christopher R.	38,800	MITRIUS, Janice V.	43,808		
HANLON, Brian E.	40,449	MORENO, Christopher P.	38,566		
HEMMENDINGER, Lisa M.	42,653	NELSON, Jon O.	24,566		
HILLMAN, Lisa	43,673	NIEGOWSKI, James A	28,331		
HONG, Patricia E.	34,373	PATEL, Binal J.	42,065		
HOSCHET, Dale H.	19,090	PATHAK, Ajay S.	38,266		

as my Attorneys.

Address all telephone calls to **Joseph P. Krause** at (312) 715-1000.

Address all correspondence to BANNER & WITCOFF, LTD., Ten South Wacker Drive, Chicago, Illinois 60606.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that wilful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor (given name, family name): Thomas J. Cloonan

Inventor's signature: Thomas J. Cloonan

Date: 7-19-00

Residence: Lisle, Illinois

Citizenship: U.S.A.

Post Office Address: 4602 Wedgewood Ct.
Lisle, Illinois 60532